arXiv:1606.00538v1 [cs.RO] 2 Jun 2016

Dictionary Learning for Robotic Grasp Recognition and Detection

Ludovic Trottier Philippe Giguère Brahim Chaib-draa
Department of Computer Science and Software Engineering
Laval University, Québec, Canada
ludovic.trottier.1@ulaval.ca
philippe.giguere@ift.ulaval.ca
brahim.chaib-draa@ift.ulaval.ca

The ability to grasp ordinary and potentially never-seen objects is an important feature in both domestic and industrial robotics. For a system to accomplish this, it must autonomously identify grasping locations by using information from various sensors, such as Microsoft Kinect 3D camera. Despite numerous progress, significant work still remains to be done in this field. To this effect, we propose a dictionary learning and sparse representation (DLSR) framework for representing RGBD images from 3D sensors in the context of determining such good grasping locations. In contrast to previously proposed approaches that relied on sophisticated regularization or very large datasets, the derived perception system has a fast training phase and can work with small datasets. It is also theoretically founded for dealing with masked-out entries, which are common with 3D sensors. We contribute by presenting a comparative study of several DLSR approach combinations for recognizing and detecting grasp candidates on the standard Cornell dataset. Importantly, experimental results show a performance improvement of 1.69% in detection and 3.16% in recognition over current state-of-the-art convolutional neural network (CNN). Even though nowadays most popular vision-based approach is CNN, this suggests that DLSR is also a viable alternative with interesting advantages that CNN has not.

1 Introduction

In robotics, automating the grasping of ordinary objects is an important open problem (Redmon and Angelova, 2015). Much progress has been made both on the hardware side (the gripper itself) and on the perception side. The development of compliant or under-actuated mechanical grippers — often referred to as "mechanical intelligence" — which passively adapt their shape to the

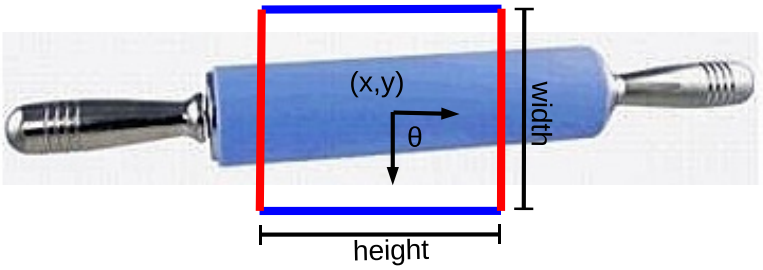

Figure 1: The grasp rectangle is a five-dimensional grasp representation. The 2D rectangle is fully determined by its center coordinates (x, y), width, height and its angle θ from the x-axis. The blue edges indicate the gripper plate location and the red edges show the gripper opening, prior to grasping.

grasped object has greatly simplified the problem (Laliberté et al., 2002). However, determining good grasping locations still requires an efficient perception system.

The advent of Microsoft Kinect inexpensive 3D camera opened the door to rapid deployment of new and robust approaches in identifying such locations. Its market accessibility and ease-of-use provided a straightforward solution for incorporating depth and RGB information (called RGBD images) into deployed systems of various industrial settings.

In this paper, we look at identifying grasping locations for two plates parallel grippers, by employing such RGBD images. As representation of a grasping location, we use Jiang *et al.* 5-dimensional *grasp rectangle* (Jiang et al., 2011) from which the 7-dimensional gripper configuration can be easily computed. The 2D oriented rectangle, shown in Fig. 1, indicates the gripper's location, orientation and physical limitations:

$$R = \{x, y, \theta, w, h\}\,.$$

Using the grasp rectangle representation makes grasp recognition analogous to object recognition (bounding box approaches), and soon for detection. For grasp recognition, the goal is to determine whether a grasp rectangle R is a good or bad candidate. For grasp detection, the goal is to predict the configuration of the best rectangle R^*. In this particular setting, identifying grasping locations can then be seen as a vision problem. This is particularly advantageous because several breakthrough works on similar vision-oriented problems have been proposed in past decades, and can thus be exploited for detecting grasping locations.

Although compelling due to its small cost and ease-of-use, the Microsoft Kinect (and most structured-light devices) has some drawbacks. One is the presence of two types of noise in depth information: the axial and lateral noise model of the object distance to the camera, and the mask noise model of missing 3D information. While vision-based approaches can reasonably deal with the former, the latter can be particularly cumbersome. Objects with shiny surfaces often cause structured-light 3D cameras to fail which results in absence of information. To cope with this phenomenon, Lenz et al. (2015) had to develop different mask-based regularization terms to solve their multi-layer neural network convergence problems caused by using zeros as masked-out entry values. With our approach, we seek to create a theoretically founded grasp localization model which can address the Kinect mask noise inherently, without resorting to a custom regularization.

The second drawback of Microsoft Kinect is the lack of available large scale RGBD grasp datasets, which makes training high-dimensional models cumbersome. As an example, Redmon and Angelova (2015) previously applied a convolutional neural network (CNN) for detecting grasp candidates that contained more than a million parameters. Due to the relative small amount of RGBD images in their grasp dataset (the Cornell dataset), they had to pre-train the CNN for several days on ImageNet (a RGB image dataset) and fine-tune for several hours on the Cornell one. In an industrial context where datasets are small and new objects are regularly added, a fast and robust training phase is essential. In particular, an efficient strategy to make objects easier to grasp is to add more images from different viewpoints and retrain, hence the importance of fast training phase.

To satisfy the aforementioned requirements, we looked at employing dictionary learning and sparse representations (DLSR). Sparse modeling of data is a biologically-inspired and theoretically founded approach in which observations are represented as linear combinations of few atoms from a dictionary. As previously shown in the context of object recognition and image restoration, DLSR is well-suited to deal with masked-out entries, has a significantly faster training phase than CNN and can work with small datasets (Wright et al., 2010). A standard DLSR method is divided into a *dictionary learning* phase, where a dictionary is trained to capture the latent structure of the data, and a *feature coding* phase, where the dictionary is used to transform raw observations into features. Representing observations by learning how to extract features from them makes DLSR particularly interesting in the actual context, as it steers clear of relying on expert knowledge brought by hand-designed feature engineering.

Our contribution with this paper is twofold. First, we propose a DLSR-based framework for learning and extracting useful information from RGBD images that is rapidly trainable, can work with a small dataset and inherently deals with masked-out entries. The goal is to demonstrate the applicability of such an approach for grasp recognition and detection, and to compare it with other ones in the literature on the standard Cornell task (shown in Fig. 2). Second, we present an empirical evaluations of several dictionary learning and feature coding approach combinations. Since DLSR has been around for some years, more than one variant exists for dictionary training and feature extraction. The large quantity of available methods makes choosing one particular combination troublesome, as few indications can guide our choice. By understanding the relationship between dictionary learning and feature coding, our goal is to ascertain which combinations are best suited for the task at hand by comparing performance, speed of training and parallelizability (either on CPU or GPU).

The rest of this paper is divided as follows. We make an overview of related works in section 2. All dictionary learning approaches and encoders are detailed in section 3, along with explanations concerning data preprocessing and the overall feature extraction process. We elaborate on the experimental framework in section 4 and report the results in section 5. Finally, we discuss the pros and cons of the approaches in section 6 and conclude in section 7.

2 Related Work

One of the fundamental concept in grasping is its representation, which has undergone significant evolution over the years. For example, Saxena et al. (2006) proposed in 2006 a 2D grasping

Figure 2: The Cornell Grasping Dataset contains images of a wide variety of everyday objects. This database defines two tasks: grasp recognition and grasp detection. For grasp recognition, the goal is to determine whether a grasp rectangle R is a good or a bad candidate. For grasp detection, the goal is to predict the configuration of the best grasp rectangle R^*.

point representation, while more recently, Le et al. (2010) proposed a pair of points. However, these representations did not faithfully represent the 7-dimensional gripper configuration and its inherent mechanical constraints, which led to the formulation of the grasp rectangle of Jiang et al. (2011) in 2011.

For identification, several previous approaches used 3D simulations to learn good grasping regions (Goldfeder et al., 2007; Miller and Allen, 2004; Detry et al., 2013; Pelossof et al., 2004). A strong limitation for them is the need to know all 3D physical models *a priori*, which severely reduces the applicability for general purpose robots. Better approaches for performing grasp identification without building complex models prior to the execution would certainly be more adequate.

Other works have shown the importance of depth information (Lai et al., 2011; Blum et al., 2012) and image processing for representing the image inputs (Maitin-Shepard et al., 2010; Saxena et al., 2008). They often rely on hand-designed features (Rusu et al., 2010), making them possibly brittle or hard to tune. While there has been some work on applying neural network for learning RGBD image features (Socher et al., 2012; Gupta et al., 2014), robotic grasping using neural networks research is still in its infancy (Lenz et al., 2015; Redmon and Angelova, 2015).

Figure 3: Left: a pair of scissors with a candidate grasp rectangle image. Center-Left: image taken from the rectangle rotated to match the global image orientation. Center-Right: rescaled image with preserved aspect ratio. The black regions indicate masked-out padding. Right: rescaled image without preserved aspect ratio. Preserving aspect ratio and using padding allows the object parts to correctly appear graspable.

The literature on Unsupervised Feature Learning (UFL) is vast. DLSR-based approaches like (Aharon et al., 2006; Mairal et al., 2009; Yang et al., 2009) achieved impressive results on vision-related tasks such as object recognition (Bo et al., 2013), face recognition (Zhang and Li, 2010), scene analysis (Lazebnik et al., 2006) and image restoration (Elad and Aharon, 2006). Conversely, not as much work has been done to evaluate the performance of dictionary learning approaches on RGBD images (Bo et al., 2013). By applying such a paradigm for identifying grasping locations, one contribution of this paper is to present a comparative study of several DLSR approach combinations which is currently lacking in grasping literature.

3 Learning Framework

In this section, we detail the dictionary learning approaches that are used for unsupervised feature learning. We also elaborate on the data preprocessing, the overall feature extraction process, the classifier for grasp recognition and the regressor for grasp detection.

3.1 Data Preprocessing

We first compute the gray channel (K) and estimate the depth normal coordinates N_x, N_y and N_z for each RGBD Kinect image, as done in Bo et al. (2013). Each image now contains eight channels. Then, we preprocess each grasp rectangle image by rotating it to match the global image orientation and by rescaling it to a $24 \times 24 \times 8$ size with aspect-ratio preserved. Fig. 3 shows an example of a grasp rectangle image that is preprocessed in such a way.

In order to learn a set of features, we first collect a batch (100,000 in our experiments) of small $6 \times 6 \times 8$ patches, extracted at random from the $24 \times 24 \times 8$ images, that are then channel-wise standardized and ZCA whitened (Hyvärinen et al., 2004). Given a set of these patch vectors $x^{(i)} \in \mathbb{R}^n$, we apply a dictionary learning approach to learn a dictionary $D \in \mathbb{R}^{n \times d}$, where each column $D^{(j)}$ is one atom that represents the latent structure of the patches. Fig. 4 shows an example of a dictionary learned with Cornell database, the same we used in our tests. Most squares (atoms) show localized and oriented Gabor-like filters that are known to be relevant features for representing raw images (Marčelja, 1980).

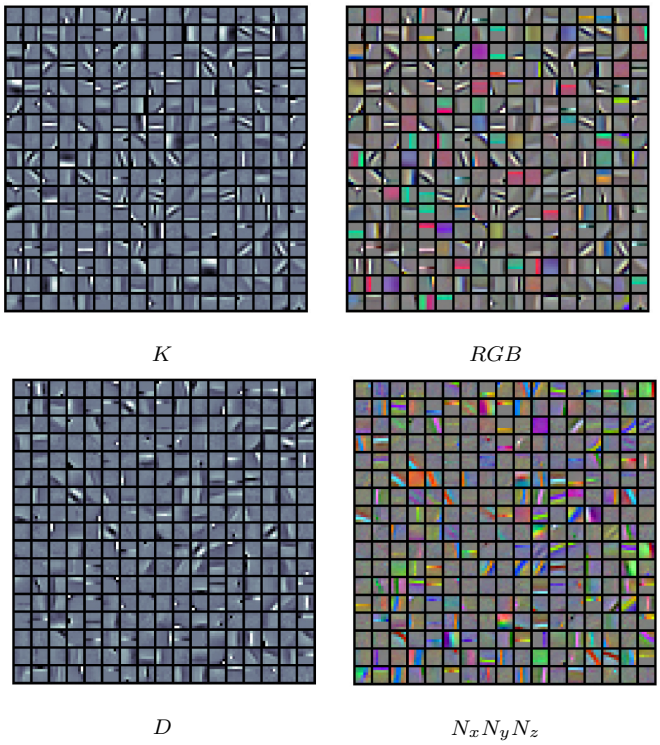

Figure 4: A dictionary D of 300 atoms (each square is an atom $D^{(j)}$) learned using the Cornell dataset shown in four distinct parts: K (gray), RGB, D (depth) and $N_x N_y N_z$ (depth normals). Most squares (atoms) show localized and oriented Gabor-like filters.

3.2 Dictionary Learning

We now elaborate on the dictionary learning algorithms that we chose for learning a dictionary D. Specifically, we tested the following approaches:

3.2.1 Sparse Coding (SC)

We train the dictionary by optimizing a ℓ_1-regularized sparse coding formulation:

$$\min_{D, w^{(i)}} \sum_i \|Dw^{(i)} - x^{(i)}\|_2^2 + \lambda \|w^{(i)}\|_1 \qquad (1)$$
$$\text{subject to } \|D^{(j)}\|_2 = 1, \, \forall j \, ,$$

using the online dictionary learning (ODL) algorithm of Mairal et al. (2009). ODL minimizes (1) alternatively over the sparse weights $w^{(i)}$ and the dictionary D, making it a fast and scalable approach for large datasets. We used Least Angle Regression (LARS) (Efron et al., 2004) to solve for the sparse weights. In our experiments, we cross-validated the sparsity parameter λ with $\lambda \in \{0.5, 1.0, 1.5, 2.0\}$.

3.2.2 Orthogonal Matching Pursuit (OMP)

In this case, the ℓ_1 penalty is replaced by a ℓ_0 one, and optimization follows a formulation similar to SC:

$$\begin{aligned} &\min_{D,w^{(i)}} \sum_i \|Dw^{(i)} - x^{(i)}\|_2^2 \\ &\text{subject to } \|D^{(j)}\|_2 = 1, \ \forall j \\ &\text{and } \|w^{(i)}\|_0 \leq \gamma, \ \forall i \, . \end{aligned} \tag{2}$$

We again used ODL (Mairal et al., 2009) to learn the dictionary, but this time solved the sparse weights with Orthogonal Matching Pursuit (OMP) (Pati et al., 1993). In our experiments, we cross-validated the sparsity parameter γ with $\gamma \in \{1, 5, 10, 15\}$.

3.2.3 Gain-Shape Vector Quantization (GSVQ)

The idea is to represent a vector by separating its gain (euclidean norm) from its shape (orientation) (Gersho and Gray, 2012). GSVQ has a similar formulation to OMP where $\gamma = 1$. Specifically, it computes $k = \arg\max_j |D^{(j)} \cdot x^{(i)}|$, the atom index that is most correlated with $x^{(i)}$, then sets $w_k^{(i)} = D^{(k)} \cdot x^{(i)}$ and $w_j^{(i)} = 0$ for $j \neq k$. Using these fixed weight vectors, it is then straightforward to find the locally optimal dictionary D in (2) using an iterative procedure as in KMeans.

3.2.4 Normalized KMeans (NKM)

Coates et al. (2011) showed that the centroids learned by a standard KMeans algorithm make good dictionary atoms. We thus clustered the patches and used the normalized centroids as dictionary atoms.

3.2.5 Randomly Sampled Patches (RP)

Here, we used the heuristic proposed by Coates and Ng (2011) to populate the dictionary. We uniformly sampled d patches from the dataset and used the normalized vectors as the dictionary atoms.

3.2.6 Random Dictionary (R)

It has been shown previously that completely random weights can achieve surprisingly good results (Coates and Ng, 2011; Saxe et al., 2011). Therefore, we also tested *learning* the dictionary by sampling d times the uniform distribution $U([0, 1]^n)$ and used the normalized vectors as dictionary atoms.

3.3 Feature Coding

Executing any dictionary learning approaches presented in the previous section gives a dictionary D representing the latent structure of the data. To actually extract features from them, we use an *encoder* that maps an observation $x^{(i)}$ to its feature representation $f^{(i)}$.

3.3.1 Sparse Coding (SC)

The first approach is based on the sparse coding formulation of (1). Given a dictionary learned with any of section 3.2 methods (not necessarily SC), we solve (1) for the sparse weights assuming a fixed D:

$$w^{(i)} = \arg\min_{w^{(i)}} \|Dw^{(i)} - x^{(i)}\|_2^2 + \lambda\|w^{(i)}\|_1, \forall i\,, \tag{3}$$

using LARS (Efron et al., 2004). It is important to note that λ (also γ and τ) may have a different value during feature coding than dictionary learning. We then apply *polarity splitting* (Coates and Ng, 2011), that is, we split the positive weights from the negative ones:

$$f_j^{(i)} = \max\{w_j^{(i)}, 0\}$$
$$f_{j+d}^{(i)} = \max\{-w_j^{(i)}, 0\}$$

This technique allows the classifier to model positive and negative weights differently, thus improving its flexibility (Coates and Ng, 2011). In our experiments, we cross-validated with $\lambda \in \{0.5, 1, 1.5, 2\}$.

3.3.2 Masked Sparse Coding (mSC)

With this approach, we explicitly deal with masked-out entries arising from either the preservation of the aspect ratio during grasp rectangle rescaling or noisy Kinect depth sensor. Let $m^{(i)} \in \{0, 1\}^n$ be the mask vector of observation $x^{(i)}$ where $m_j^{(i)} = 0$ implies that $x_j^{(i)} = 0$ indicates a masked entry (similarly, $m_j^{(i)} = 1$ indicates that $x_j^{(i)} \gtreqless 0$ is not a masked entry). We then remove the penalty induced by the masked entries in (3) giving the following formulation, $\forall i$:

$$w^{(i)} = \arg\min_{w^{(i)}} \|\text{diag}(m)(Dw^{(i)} - x^{(i)})\|_2^2 + \lambda\|w^{(i)}\|_1 \tag{4}$$

To solve (4), we again used LARS (Efron et al., 2004), this time with the mask. As in SC, we apply polarity splitting. In our experiments, we cross-validated with $\lambda \in \{1, 2, 3, 4\}$.

3.3.3 Orthogonal Matching Pursuit (OMP)

This approach is based on the formulation of (2). Assuming a fixed dictionary, we applied OMP (Pati et al., 1993) to solve for the sparse weights:

$$w^{(i)} = \min_{w^{(i)}} \|Dw^{(i)} - x^{(i)}\|_2^2 \tag{5}$$
$$\text{subject to } \|w^{(i)}\|_0 \leq \gamma, \forall i\,.$$

We again used polarity splitting as in SC. We cross-validated with $\gamma \in \{1, 5, 10, 15\}$.

3.3.4 Masked Orthogonal Matching Pursuit (mOMP)

Similarly to mSC, we remove the penalty of the masked entries from (5) giving the following formulation

$$w^{(i)} = \min_{w^{(i)}} \|\text{diag}(m)(Dw^{(i)} - x^{(i)})\|_2^2 \quad (6)$$
$$\text{subject to } \|w^{(i)}\|_0 \leq \gamma, \forall i\,,$$

which we solved using OMP (Pati et al., 1993) but this time with the mask. We used polarity splitting and cross-validated with $\gamma \in \{1, 5, 10, 15\}$.

3.3.5 Soft-Thresholding (ST)

Soft-Thresholding (Donoho and Johnstone, 1995) (also known as marginal regression (Genovese et al., 2012)) is a fast alternative to finding the optimal solution of (3). It is based on the (strong) hypothesis that all weights $w_j^{(i)}$ are independent. This enables solving (3) for the sparse weights marginally (each $w_j^{(i)}$ individually) thus giving a simple analytical solution:

$$w_j^{(i)} = \text{sign}(D^{(j)} \cdot x^{(i)}) \cdot \max\left\{0, |D^{(j)} \cdot x^{(i)}| - \tau\right\}\,, \quad (7)$$

where τ is the sparsity parameter. Similar to SC and OMP, we applied polarity splitting on $w^{(i)}$. We cross-validated with $\tau \in \{0.5, 1, 1.5, 2\}$.

3.3.6 Natural (N)

Finally, we define a *natural* encoder as whichever approach was used for solving the sparse weights during dictionary learning. For instance, the natural encoder of SC dictionary learning is SC feature coding with the same sparsity parameter λ. Similarly, the natural encoder of OMP dictionary learning is OMP feature coding with the same γ. We used a different approach for the other dictionary learning algorithms, since they do not require a sparsity parameter. Specifically, for GSVQ we used OMP with $\gamma = 1$. For R and RP we used ST with $\tau = 0$ which corresponds to a random linear projection. Finally, for NKM we did not normalize the centroids (as in standard KMeans) and used the KMeans-Tri feature coding of Coates et al. (2011).

3.4 Feature Extraction Process

Given a learned dictionary D and a choice of encoder, a patch vector $x^{(i)}$ can now be transformed into its feature representation $f^{(i)}$. Here we describe how to transform a full grasp rectangle image into a feature representation usable by the classifier for grasp recognition. Our feature extraction process follows the spatial pyramid matching framework proposed by Yang et al. (2009). The entire process is shown in Fig. 5. Specifically, we extract a batch of $6 \times 6 \times 8$ patches $x^{(i)}$ in a convolutional way with a stride of one pixel. These patches are channel-wise standardized and ZCA whitened. Then, we map each patch to its feature representation $f^{(i)}$ using the dictionary D and the encoder. After, we divide the image into four quadrants and perform sum pooling over the feature vectors $f^{(i)}$ of each quadrant respectively. Finally, the pooled feature vectors from all quadrants are concatenated into a single vector that is used as input to the classifier.

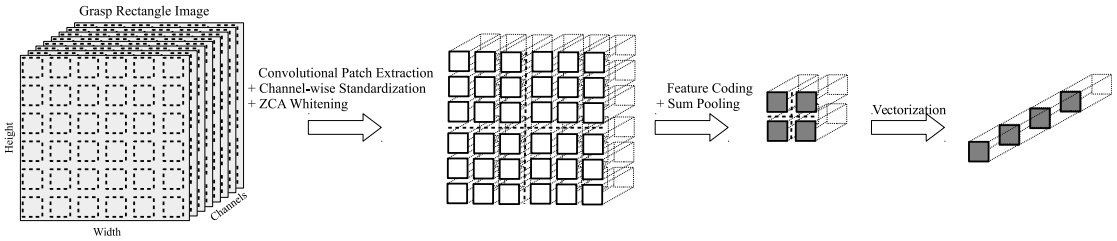

Figure 5: Feature extraction process. Left: a batch of $6 \times 6 \times 8$ patches $x^{(i)}$ are extracted in a convolutional way (with a stride of one pixel) from the image. The patches are then channel respective standardized and ZCA whitened. Center-Left: each patch $x^{(i)}$ is mapped to its feature representation $f^{(i)}$ given a dictionary D and a choice of encoder. Center-Right: a four quadrants sum pooling is applied on feature vectors $f^{(i)}$. Right: all quadrant pooled weights are concatenated into a single vector. The final vector on the right is used as input to the SVM.

3.5 Grasp Recognition Classifier

For classification, we optimized a ℓ_2-linear SVM using a standard L-BFGS solver from Schmidt's *minFunc* toolbox (Schmidt, 2005). We cross-validated the regularization parameter C with $C \in \{1, 10, 100, 1000\}$.

3.6 Grasp Detection Regressor

Although performing grasp recognition is straightforward with a SVM because it is a binary classification problem, grasp detection is more cumbersome. Directly predicting the best 5-dimensional grasp rectangle from the very high-dimensional inputs ($640 \times 480 \times 4$ Kinect images) would not be realistic with our current framework. A naive application of the feature extraction process described in section 3.4 on an entire Kinect image would extract an unreasonably high-dimensional feature vector with little discriminative power. We instead opted for a standard grid-search in grasp rectangle space. We first performed background removal and identified the smallest region containing the object. Then, we extracted, in a convolutional way with a stride of 10 pixels, grasp rectangles from the image with varying sizes and orientations. We varied the size of the rectangle from 10 pixels to 90 pixels with a stride of 10 pixels, and varied the orientation from 0 degree to 180 degrees with a stride of 15 degrees. For each of these grasp rectangle images, we performed feature extraction and inputed them to the SVM. The rectangle having the highest classification score was chosen as the candidate grasp.

4 Experimental Framework

4.1 Cornell Dataset

The Cornell Grasping Dataset (Jiang et al., 2011) contains 885 RGBD images of 240 distinct objects (available at `http://pr.cs.cornell.edu/deepgrasping/`). Each image has multiple

positively- and negatively-labeled grasp rectangles, specifically selected for parallel plate grippers. The labeled rectangles are varied in terms of size, orientation and position, but are by no means exhaustive of every grasp scenarios (some image have graspable regions that are not labeled).

4.2 Grasp Recognition Experiments

For grasp recognition, we performed the following three experiments:

4.2.1 Grasp Recognition Evaluation

Previous works on the Cornell dataset reported their results using a 5-fold cross-validation (Jiang et al., 2011; Lenz et al., 2015). They optimized for the hyper-parameters using a separate set of grasp examples (which we call the validation set). However, exactly comparing our results to theirs is impossible because they did not report which examples they selected for validation. Therefore, we instead report our recognition accuracies using a 5-5 folds nested cross-validation. The advantages are that this removes the need to specify the validation set and reduces the bias induced by choosing which examples to put in it. Moreover, our results are still comparable to the previous ones and, most importantly, allows future works on the dataset to report results based on the same evaluation framework.

4.2.2 Varying the Dictionary Size

The number d of atoms has a direct influence on the performance of the classifier. On one side, a too small dictionary D does not capture enough structure to correctly represent the raw data. On the contrary, a too big dictionary D contains noisy atoms which are never activated in the sparse weight vectors $w^{(i)}$. These unused atoms slow down the weight extraction process and make the resulting feature vector $f^{(i)}$ unnecessarily long and noisy. We therefore evaluate the correlation between the recognition accuracy and the dictionary size d to later guide our choice of it during grasp detection.

4.2.3 ZCA Whitening

Previous works on dictionary learning applied to RGB object recognition have already shown that whitening improves the accuracy of the classifier (Coates et al., 2011; Coates and Ng, 2011). We therefore wanted to validate that it is still the case when applying DLSR approaches for grasp recognition in RGBD images.

4.3 Grasp Detection Experiments

For grasp detection, we performed the following two experiments:

4.3.1 Grasp Detection Evaluation

To evaluate the quality of a grasp candidate, we used the *rectangle metric* as in (Jiang et al., 2011; Lenz et al., 2015). Specifically, if the rectangle metric of any of the ground truth rectangles with

the candidate is positive, the regression is a success. In more detail, the metric is positive if: *1)* the candidate orientation is within 30° of the ground truth rectangle, and *2)* the Jaccard index between the candidate and the ground truth is greater than 25%, where the Jaccard index between two rectangles R_1 and R_2 is defined as:

$$J\left(R_1, R_2\right) = \frac{\text{area}(R_1 \cap R_2)}{\text{area}(R_1 \cup R_2)}\,. \tag{8}$$

Unlike grasp recognition, we performed a standard 5 folds cross-validation for the detection problem. We did not optimize for the hyper-parameters, but instead used those that were the most often selected in the nested (second layer) cross-validation during grasp recognition. We also used two learning scenarios (Lenz et al., 2015):

- Image-wise splitting: where we split the images randomly.
- Object-wise splitting: where we split the objects randomly, gathering all the image of the object in the same fold.

Image-wise splitting studies the ability to generalize to new positions and orientations of an object that has already been seen. Object-wise splitting examine the capability to generalize to novel, unseen objects. While the first scenario is more suitable in an industrial context, because the set of objects is known beforehand, the second one is more difficult but also more realistic. Since training on all possible objects is almost impossible, this help asserting the viability of the approach to perform everyday grasping.

4.3.2 Self-Taught Learning

One of the most appealing advantage of unsupervised feature learning is the possibility of using unlabeled data for learning the dictionary D. We therefore evaluated the approaches on the problem of *self-taught learning* (Raina et al., 2007). Specifically, we randomly subsampled 30 objects from the Washington RGBD dataset, which also contains Kinect images of everyday objects (Lai et al., 2011). We then randomly selected 25 images per object, giving 750 images in total. From these images, we further extracted 100,000 $6 \times 6 \times 8$ patches and added them to the 100,000 patches extracted from the Cornell dataset images. We then learned the dictionary D using the patches from both datasets. This test examined the capacity of learning useful features from images taken from another distribution that will never be tested on.

5 Experimental Results

5.1 Grasp Recognition Results

5.1.1 Grasp Recognition Evaluation

The nested cross-validation accuracies (in $\%$) for grasp recognition using a dictionary of d=300 atoms are reported in Table 1. The best dictionary learning + encoder combination was NKM-SC, which reached an accuracy of 96.86%, while the lowest accuracy is obtained with R-mOMP at

Table 1: Cross-validation results of all combinations of dictionary learning and feature coding, for Cornell dataset. Numbers are grasp recognition accuracies, in percent (%), from 5-5 folds nested cross validation where hyper-parameter maximization is performed on the nested folds.

Learning \ Features	SC	mSC	OMP	mOMP	ST	Natural
SC	96.63	96.74	96.61	96.61	96.73	96.65
OMP	96.68	96.71	96.60	96.69	96.50	96.58
GSVQ	96.72	96.50	96.66	96.70	96.50	95.65
NKM	96.86	96.64	96.58	96.64	96.42	96.52
RP	96.43	96.51	96.35	96.20	96.28	96.37
R	95.84	95.52	95.34	95.15	95.78	95.90

95.15%. As a comparison, previous approaches from Jiang et al. (2011), who used a cascade of multi-layer perceptrons, achieved 89.6% and Lenz et al. (2015), who used a ImageNet pre-trained convolutional neural network (CNN), had 93.7%. These results suggests that even though neural network-based models are nowadays most popular way to solve vision-related problems, DLSR is still a viable approach because *1)* it obtained the highest accuracy on Cornell recognition task and *2)* the training cost is significantly smaller than CNN in both the computation time and the training dataset size.

Table 1 shows that, apart from method R, any DLSR combination arrives at similar performances (accuracies vary by no more than 1%). This suggests that selecting a dictionary learning or a feature coding approach may only be based on the practical consideration that it requires no hyper-parameter tuning. For instance, GSVQ, NKM and RP dictionary learning approaches could be considered before SC and OMP due to their hyper-parameter free nature (SC has λ and OMP has γ). Similarly, NKM, GSVQ and RP natural features may be taken in consideration before all other feature encoding because they have no hyper-parameters.

The best overall encoders across the dictionary learning approaches are SC and mSC. The ℓ_1 optimization (3), along with its masked version (4), appear to extract the most robust features. However, these encoders were the most time consuming of all, due to the tedious optimization needed for solving the sparse weights $w^{(i)}$. For real-time scenarios, these encoders would be too cumbersome to use. We thus select GSVQ-ST, NKM-Natural and RP-Natural as the three most appealing combinations and use them for the next experiments. NKM-Natural and RP-Natural are both hyper-parameter free, have a fast encoder and dictionary learning algorithm. Even though GSVQ-ST defines the λ hyper-parameter, it is a fast encoder which shows good recognition accuracies. For more thorough evaluations, we also keep the top performer NKM-SC and OMP-Natural for its fast greedy optimization.

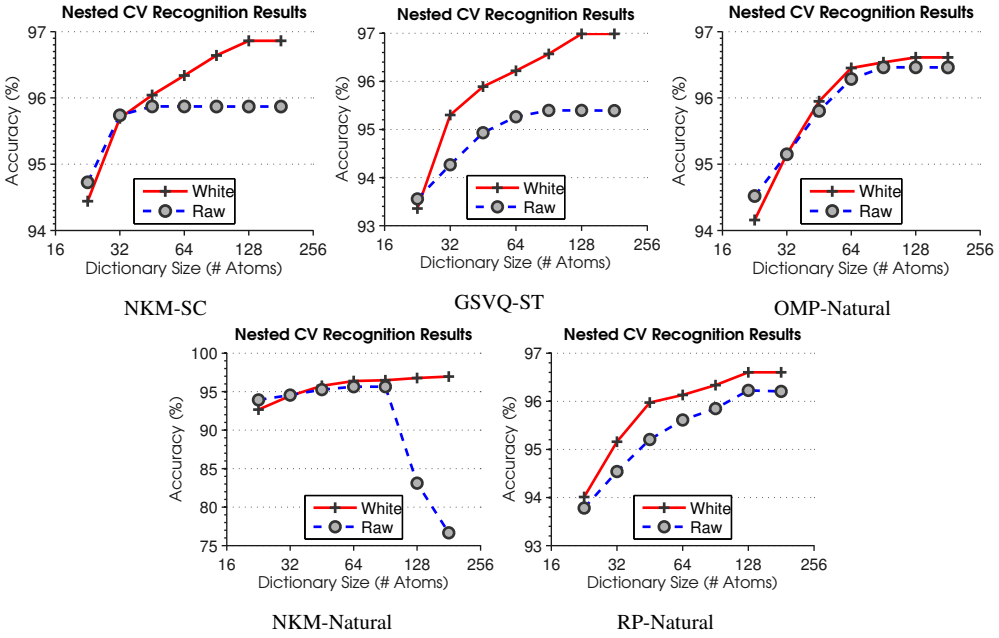

Figure 6: The effects of whitening and the dictionary size on recognition accuracies. Whitening improves the performance and may degrade it when not used. Increasing the dictionary size improves the results up to a limit (plateau) where no more boost is possible.

5.1.2 The Effects of Whitening and Dictionary Size

The effects of whitening along with varying the dictionary size d are displayed in Fig. 6. Whitening improves performance, and can degrade results when not used, as seen with NKM-Natural. Increasing the dictionary size d improves the results up to a limit where no additional gain is possible. The plateau indicates that dictionary learning is unable to learn new useful features, thus reaching a limit in its representative capability. We can see visually the impact of whitening in Fig. 7. The dictionary learned with raw images clearly shows a lack of Gabor-like filters. This is because whitening removes redundant information in the inputs, hence learning more discriminative features.

5.2 Grasp Detection Results

5.2.1 Grasp Detection Evaluation

The cross-validation accuracies (in %) for grasp detection of the five selected approaches using a dictionary of $d = 300$ atoms are reported in Table 2. For image-wise and object split respectively, the best accuracies are obtained by NKM-Natural with 89.40% and GSVQ-ST with 88.79%. As comparison, Lenz et al. (2015) obtained 73.9% and 75.6% with a cascade of multi-layer perceptrons while Redmon and Angelova (2015) achieved 88.0% and 87.1% with a CNN.

Even though CNN is a powerful approach (currently state-of-the-art in several vision-related problems), DLSR has one advantage CNN has not. Due to the relative small amount of data

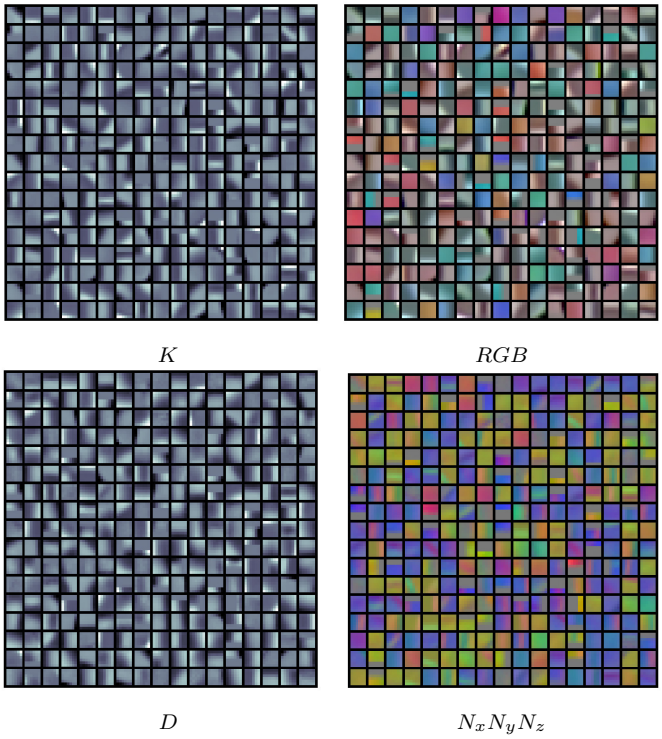

Figure 7: A dictionary D of 300 atoms (each square is an atom $D^{(j)}$) learned using the Cornell dataset shown in four distinct parts: K (gray), RGB, D (depth) and $N_x N_y N_z$ (depth normals). No whitening is performed, and we clearly see the absence of localized and oriented Gabor-like filters.

in Cornell dataset, Redmon *et al.* pre-trained their CNN on ImageNet containing RGB images, replaced blue with depth channel to make it compatible with RGBD images (giving RGD images), and performed a final fine-tuning on Cornell images. Since blue channel-related low level features are unlikely extracting useful information from depth, such a pre-training approach is clearly sub-optimal. Directly training the CNN on Cornell dataset would require gathering a large quantity of additional images, at a substantial cost. In contrast, DLSR can be directly trained on Cornell RGBD images despite its small size, and this makes it more advantageous in this manner.

5.2.2 Self-Taught Learning

The cross-validation accuracies (in %) for grasp detection of the five selected approaches using a dictionary of 300 atoms and patches from both Cornell and Washington datasets are reported in Table 3. Even tough we see both a performance improvement (NKM-SC and NKM-Natural) and decrease (GSVQ-ST, OMP-Natural and RP-Natural), the variations are small (around 0.5%) and not significant. This suggests that the dictionary learning approaches were able to extract all the necessary features from the Cornell images, and the additional features extracted from the Washington dataset were not helpful. This is further reflected in Fig. 4 by the presence of some

Table 2: Cross validation detection results for the Cornell dataset. See text for a discussion on computing time.

Algorithm	Detection Accuracy (%)	
	Image-wise Split	Object-wise Split
Jiang et al. (2011)	60.5	58.3
Lenz et al. (2015)	73.9	75.6
Redmon and Angelova (2015)	88.0	87.1
NKM-SC	88.67	88.07
GSVQ-ST	88.72	**88.79**
OMP-Natural	89.34	88.56
NKM-Natural	**89.40**	88.17
RP-Natural	87.70	86.61

Table 3: Cross validation detection results for the Cornell dataset using patches from both Cornell and Washington datasets.

Algorithm	Detection Accuracy (%)	
	Standard	Self-Taught
NKM-SC	88.07	88.85
GSVQ-ST	88.79	87.76
OMP-Natural	88.56	88.18
NKM-Natural	88.17	88.53
RP-Natural	86.61	86.12

weakly structured atoms (blank squares). Indeed, a dictionary learning approach normally use all available atoms to well represent highly structured data. Here, it could achieve that by using only a subset of the atoms, thus showing that it did not need additional patches to learned all the relevant features.

6 Discussion

To obtain Table 2 high accuracy results in detection, we had to pay a computational price. Even though feature coding approaches have reasonable low computational complexity (SC is $O(d^3 + nd^2)$, OMP is $O(\gamma nd)$ and ST is $O(nd)$), the exhaustive grid search in grasp rectangle space is computationally demanding. This translates into several minutes to complete a detection which is higher than Redmon *et al.*'s CNN with 76 ms per image. However, since we used a standard CPU and they used a high-end GPU, a DLSR GPU implementation would make a fairer time computation comparison. While implementing ST would be simple, this is not

straightforward for SC and OMP due to their recursive nature. A possible avenue for a useful GPU implementation may be to parallelize grid search by exploiting grasp rectangle candidate independence, i.e. by extracting and scoring all candidates in parallel. However, grid search in grasp rectangle space being more cumbersome than spatial convolutions, the computational time of such a parallelization would still be higher than CNN.

While DLSR is fairly slow during detection, the training phase is significantly faster than CNN. Training a CNN (as the one used by Redmon and Angelova (2015)) take several days with parallel high-end GPUs, and fine-tuning on Cornell dataset takes several hours. In comparison, it took approximatively ten minutes to train our 300 atoms dictionaries on a standard CPU. While fast training phase is irrelevant in real-time test scenarios, it could be useful to train a CNN directly on RGBD images. Redmon and Angelova (2015) pre-training on ImageNet could be avoided by greedily stacking dictionaries learned on RGBD images, as previously proposed by Bengio et al. (2007) with auto-encoders in the context of RGB images. Such a DLSR and CNN combination would bring the best of both approaches, in which DLSR would improve training while CNN would provide fast detection. We intend to investigate this avenue in future works.

Even though it achieved the lowest detection accuracies, the RP-Natural combination is appealing because training the dictionary is instantaneous, feature coding requires only a matrix multiplication, and there is no hyper-parameters. Due to its simplicity, integrating the approach to a grasp localization system in its early deployment phase is straightforward and can give a good glimpse of the overall system performance in later deployment phases. One interesting avenue for future work is to understand the reason why input decorrelation allows randomly sampled patches to make such good dictionaries. For instance, is linear independence sufficient, or better dictionaries could be recovered with non-linear independence? These are several avenues worth investigating.

7 Conclusion

A perception system that determines good grasping positions from Microsoft Kinect RGBD images is a key element toward automating the grasping of ordinary objects. In this paper, we proposed a DLSR framework to recognize and detect grasp rectangles on images of object to be held by two-plates parallel grippers. Our comparative study of various dictionary learning and feature coding approach combinations on Cornell dataset have shown that the proposed DLSR framework outperformed previous neural network-based approaches. As opposed to CNN, the best DLSR combination obtained a greater accuracy in both grasp recognition and detection task despite training only on small amount of images. In addition to having a substantially fast training phase, DLSR can inherently deal with masked-out entries in noisy depth maps and do not rely on sophisticated regularization terms. As discussed in section 6, exploiting DLSR fast training phase may be a suitable research avenue for future work. Stacking dictionaries learned on RGBD images to pre-train a CNN would bring the best of both approaches, in which DLSR improve training while CNN provide fast detection.

References

Aharon, M., Elad, M., and Bruckstein, A. (2006). KSVD: An algorithm for designing overcomplete dictionaries for sparse sepresentation. *Signal Processing, Transactions on*, 54(11):4311–4322.

Bengio, Y., Lamblin, P., Popovici, D., Larochelle, H., et al. (2007). Greedy layer-wise training of deep networks. *Advances in neural information processing systems*, 19:153.

Blum, M., Springenberg, J. T., Wülfing, J., and Riedmiller, M. (2012). A learned feature descriptor for object recognition in RGB-D data. In *ICRA*, pages 1298–1303.

Bo, L., Ren, X., and Fox, D. (2013). Unsupervised feature learning for RGB-D based object recognition. In *Experimental Robotics*, pages 387–402.

Coates, A. and Ng, A. Y. (2011). The importance of encoding versus training with sparse coding and vector quantization. In *ICML*, pages 921–928.

Coates, A., Ng, A. Y., and Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, pages 215–223.

Detry, R., Ek, C. H., Madry, M., and Kragic, D. (2013). Learning a dictionary of prototypical grasp-predicting parts from grasping experience. In *ICRA*, pages 601–608.

Donoho, D. L. and Johnstone, I. M. (1995). Adapting to unknown smoothness via wavelet shrinkage. *Journal of the american statistical association*, 90(432):1200–1224.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *The Annals of statistics*, 32(2):407–499.

Elad, M. and Aharon, M. (2006). Image denoising via sparse and redundant representations over learned dictionaries. *Image Processing, Transactions on*, 15(12):3736–3745.

Genovese, C. R., Jin, J., Wasserman, L., and Yao, Z. (2012). A comparison of the lasso and marginal regression. *JMLR*, 13(1):2107–2143.

Gersho, A. and Gray, R. M. (2012). *Vector quantization and signal compression*, volume 159. Springer Science & Business Media.

Goldfeder, C., Allen, P. K., Lackner, C., and Pelossof, R. (2007). Grasp planning via decomposition trees. In *ICRA*, pages 4679–4684.

Gupta, S., Girshick, R., Arbeláez, P., and Malik, J. (2014). Learning rich features from RGB-D images for object detection and segmentation. In *ECCV*, pages 345–360.

Hyvärinen, A., Karhunen, J., and Oja, E. (2004). *Independent component analysis*, volume 46. John Wiley & Sons.

Jiang, Y., Moseson, S., and Saxena, A. (2011). Efficient grasping from RGBD images: Learning using a new rectangle representation. In *ICRA*, pages 3304–3311.

Lai, K., Bo, L., Ren, X., and Fox, D. (2011). A large-scale hierarchical multi-view RGB-D object dataset. In *ICRA*, pages 1817–1824.

Laliberté, T., Birglen, L., and Gosselin, C. (2002). Underactuation in robotic grasping hands. *Machine Intelligence & Robotic Control*, 4(3):1–11.

Lazebnik, S., Schmid, C., and Ponce, J. (2006). Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, volume 2, pages 2169–2178.

Le, Q. V., Kamm, D., Kara, A. F., and Ng, A. Y. (2010). Learning to grasp objects with multiple contact points. In *ICRA*, pages 5062–5069.

Lenz, I., Lee, H., and Saxena, A. (2015). Deep learning for detecting robotic grasps. *IJRR*, 34(4-5):705–724.

Mairal, J., Bach, F., Ponce, J., and Sapiro, G. (2009). Online dictionary learning for sparse coding. In *ICML*, pages 689–696.

Maitin-Shepard, J., Cusumano-Towner, M., Lei, J., and Abbeel, P. (2010). Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *ICRA*, pages 2308–2315.

Marčelja, S. (1980). Mathematical description of the responses of simple cortical cells. *JOSA*, 70(11):1297–1300.

Miller, A. T. and Allen, P. K. (2004). Graspit!: A versatile simulator for robotic grasping. *Robotics & Automation Magazine*, 11(4):110–122.

Pati, Y. C., Rezaiifar, R., and Krishnaprasad, P. (1993). Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Signals, Systems and Computers. Twenty-Seventh Asilomar Conference on*, pages 40–44.

Pelossof, R., Miller, A., Allen, P., and Jebara, T. (2004). An SVM learning approach to robotic grasping. In *ICRA*, volume 4, pages 3512–3518.

Raina, R., Battle, A., Lee, H., Packer, B., and Ng, A. Y. (2007). Self-taught learning: Transfer learning from unlabeled data. In *ICML*, pages 759–766.

Redmon, J. and Angelova, A. (2015). Real-time grasp detection using convolutional neural networks. In *ICRA*, pages 1316–1322.

Rusu, R. B., Bradski, G., Thibaux, R., and Hsu, J. (2010). Fast 3D recognition and pose using the viewpoint feature histogram. In *IROS*, pages 2155–2162.

Saxe, A., Koh, P. W., Chen, Z., Bhand, M., Suresh, B., and Ng, A. Y. (2011). On random weights and unsupervised feature learning. In *ICML*, pages 1089–1096.

Saxena, A., Driemeyer, J., Kearns, J., and Ng, A. Y. (2006). Robotic grasping of novel objects. In *NIPS*, pages 1209–1216.

Saxena, A., Driemeyer, J., and Ng, A. Y. (2008). Robotic grasping of novel objects using vision. *IJRR*, 27(2):157–173.

Schmidt, M. (2005). minfunc: unconstrained differentiable multivariate optimization in matlab.

Socher, R., Huval, B., Bath, B., Manning, C. D., and Ng, A. Y. (2012). Convolutional-recursive deep learning for 3D object classification. In *NIPS*, pages 665–673.

Wright, J., Ma, Y., Mairal, J., Sapiro, G., Huang, T. S., and Yan, S. (2010). Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE*, 98(6):1031–1044.

Yang, J., Yu, K., Gong, Y., and Huang, T. (2009). Linear spatial pyramid matching using sparse coding for image classification. In *CVPR*, pages 1794–1801.

Zhang, Q. and Li, B. (2010). Discriminative K-SVD for dictionary learning in face recognition. In *CVPR*, pages 2691–2698.